

Learning Residual Finite-State Automata Using Observation Tables

Anna Kasprzik

FB IV, University of Trier

kasprzik@informatik.uni-trier.de

We define a two-step learner for RFSAs based on an observation table by using an algorithm for minimal DFAs to build a table for the reversal of the language in question and showing that we can derive the minimal RFSA from it after some simple modifications. We compare the algorithm to two other table-based ones of which one (by Bollig et al. [8]) infers a RFSA directly, and the other is another two-step learner proposed by the author. We focus on the criterion of query complexity.

Keywords: Grammatical inference, residual languages, observation tables

1 Introduction

The area of grammatical inference tackles the problem of inferring a description of a formal language (a grammar, an automaton) from given examples or other kinds of information sources. Various settings have been formulated and quite a lot of learning algorithms have been developed for them. One of the best studied classes with respect to algorithmical learnability is the class of *regular languages*.

A significant part of these algorithms, of which Angluin's L^* [1] was one of the first, use the concept of an *observation table*. If a table fulfils certain conditions we can directly derive a deterministic finite-state automaton (DFA) from it, and if the information suffices this is the minimal DFA for the language in question.

In the worst case the minimal DFA has exponentially more states than a minimal NFA for a language L , and as for many applications a small number of states is desirable it seems worth to consider if we cannot obtain an NFA instead. Denis et al. [4] introduce special NFAs – *residual finite-state automata* (RFSAs) – where each state represents a residual language of L . Every regular language has a unique minimal RFSA. Denis et al. give several learning algorithms for RFSAs [5, 6, 7], which, however, all work by adding or deleting states in an automaton.

We define a two-step learner for RFSAs based on an observation table by using an algorithm for minimal DFAs to build a table with certain properties for the reversal of the language L and showing that we can derive the minimal RFSA for L from this table after some simple modifications. We compare the algorithm to two other table-based ones of which one is an incremental Angluin-style algorithm by Bollig et al. [8] which infers a RFSA directly, and the other is another two-step algorithm proposed below. The comparison mainly focuses on query complexity. We find that in theory the algorithm in [8] does not outperform the combination of known algorithms inferring the minimal DFA with the modifications we propose (although it is shown in [8] that their algorithm behaves better in practice).

2 Basic notions and definitions

Definition 1 An observation table is a triple $T = (S, E, \text{obs})$ with $S, E \subseteq \Sigma^*$ finite, non-empty for some alphabet Σ and $\text{obs} : S \times E \rightarrow \{0, 1\}$ a function with $\text{obs}(s, e) = 1$ if $se \in L$, and $\text{obs}(s, e) = 0$ if $se \notin L$.

The row of $s \in S$ is $\text{row}(s) := \{(e, \text{obs}(s, e)) | e \in E\}$, and the column of $e \in E$ is $\text{col}(e) := \{(s, \text{obs}(s, e)) | s \in S\}$. S is partitioned into two sets RED and BLUE where $uv \in \text{RED} \Rightarrow u \in \text{RED}$ for $u, v \in \Sigma^*$ (prefix-closedness), and $\text{BLUE} := \{sa \in S \setminus \text{RED} | s \in \text{RED}, a \in \Sigma\}$.

Definition 2 Let $T = (S, E, \text{obs})$ with $S = \text{RED} \cup \text{BLUE}$. Two elements $r, s \in S$ are obviously different (denoted by $r <> s$) iff $\exists e \in E$ such that $\text{obs}(r, e) \neq \text{obs}(s, e)$. T is closed iff $\neg \exists s \in \text{BLUE} : \forall r \in \text{RED} : r <> s$. T is consistent iff $\forall s_1, s_2 \in \text{RED}, s_1a, s_2a \in S, a \in \Sigma : \text{row}(s_1) = \text{row}(s_2) \Rightarrow \text{row}(s_1a) = \text{row}(s_2a)$.

Definition 3 A finite-state automaton is a tuple $\mathcal{A} = (\Sigma, Q, Q_0, F, \delta)$ with finite input alphabet Σ , finite non-empty state set Q , set of start states $Q_0 \subseteq Q$, set of final accepting states $F \subseteq Q$, and a transition function $\delta : Q \times \Sigma \rightarrow 2^Q$.

If $Q_0 = \{q_0\}$ and δ maps at most one state to any pair in $Q \times \Sigma$ the automaton is deterministic (a DFA), otherwise non-deterministic (an NFA). If δ maps at least one state to every pair in $Q \times \Sigma$ the automaton is total, otherwise partial.

The transition function can always be extended to $\delta : Q \times \Sigma^* \rightarrow 2^Q$ defined by $\delta(q, \varepsilon) = \{q\}$ and $\delta(q, wa) = \delta(\delta(q, w), a)$ for $q \in Q, a \in \Sigma$, and $w \in \Sigma^*$.

Let $\delta(Q', w) := \bigcup \{\delta(q, w) | q \in Q'\}$ for $Q' \subseteq Q$ and $w \in \Sigma^*$. A state $q \in Q$ is reachable if there is $w \in \Sigma^*$ with $q \in \delta(Q_0, w)$. A state $q \in Q$ is useful if there are $w_1, w_2 \in \Sigma^*$ with $q \in \delta(Q_0, w_1)$ and $\delta(q, w_2) \cap F \neq \emptyset$, otherwise useless.

The language accepted by \mathcal{A} is $\mathcal{L}(\mathcal{A}) := \{w \in \Sigma^* | \delta(Q_0, w) \cap F \neq \emptyset\}$.

From $T = (S, E, \text{obs})$ with $S = \text{RED} \cup \text{BLUE}$ and $\varepsilon \in E$ derive an automaton $\mathcal{A}_T := (\Sigma, Q_T, Q_{T0}, F_T, \delta_T)$ defined by $Q_T = \text{row}(\text{RED})$, $Q_{T0} = \{\text{row}(\varepsilon)\}$, $F_T = \{\text{row}(s) | \text{obs}(s, \varepsilon) = 1, s \in \text{RED}\}$, and $\delta_T(\text{row}(s), a) = \{q \in Q_T | \neg(q <> \text{row}(sa)), s \in \text{RED}, a \in \Sigma, sa \in S\}$. \mathcal{A}_T is a DFA iff T is consistent. The DFA for a regular language L derived from a closed and consistent table has the minimal number of states (see [1], Th. 1). This DFA is the *canonical DFA* \mathcal{A}_L for L and is unique.

The Myhill-Nerode equivalence relation \equiv_L is defined by $r \equiv_L s$ iff $re \in L \Leftrightarrow se \in L$ for all $r, s, e \in \Sigma^*$. The index of L is $I_L := |\{[s_0]_L | s_0 \in \Sigma^*\}|$ where $[s_0]_L$ is the equivalence class under \equiv_L containing s_0 .

Theorem 1 (Myhill-Nerode theorem) – see for example [3])

I_L is finite $\Leftrightarrow L$ can be recognized by a finite-state automaton $\Leftrightarrow L$ is regular.

\mathcal{A}_L has exactly I_L states, each of which represents an equivalence class under \equiv_L .

Definition 4 The reversal \overline{w} of $w \in \Sigma^*$ is defined inductively by $\overline{\varepsilon} := \varepsilon$ and $\overline{wa} := \overline{w}a$ for $a \in \Sigma, w \in \Sigma^*$. The reversal of $X \subseteq \Sigma^*$ is defined as $\overline{X} := \{\overline{w} | w \in X\}$. The reversal of an automaton $\mathcal{A} = (\Sigma, Q, Q_0, F, \delta)$ is defined as $\overline{\mathcal{A}} := (\Sigma, Q, F, Q_0, \overline{\delta})$ with $\overline{\delta}(q', w) = \{q \in Q | q' \in \delta(q, w)\}$ for $q' \in Q, w \in \Sigma^*$.

Definition 5 The residual language (RL) of $L \subseteq \Sigma^*$ with regard to $w \in \Sigma^*$ is defined as $w^{-1}L := \{v \in \Sigma^* | wv \in L\}$. A RL $w^{-1}L$ is called prime iff $\bigcup \{v^{-1}L | v^{-1}L \subsetneq w^{-1}L\} \subsetneq w^{-1}L$, otherwise composed.

By Theorem 1 the set of distinct RLs of a language L is finite iff L is regular. There is a bijection between the RLs of L and the states of the minimal DFA $\mathcal{A}_L = (\Sigma, Q_L, \{q_L\}, F_L, \delta_L)$ defined by $\{w^{-1}L \mapsto q' | w \in \Sigma^*, \delta_L(q_L, w) = \{q'\}\}$.

Let $L_q := \{w | \delta(q, w) \cap F \neq \emptyset\}$ for a regular language $L \subseteq \Sigma^*$, some automaton $\mathcal{A} = (\Sigma, Q, Q_0, F, \delta)$ recognizing L , and $q \in Q$.

Definition 6 A residual finite-state automaton (RFSA) is an NFA $\mathcal{A} = (\Sigma, Q, Q_0, F, \delta)$ such that L_q is a RL of $\mathcal{L}(\mathcal{A})$ for all states $q \in Q$.

Definition 7 The canonical RFSA $\mathcal{R}_L = (\Sigma, Q_R, Q_{R0}, F_R, \delta_R)$ for $L \subseteq \Sigma^*$ is defined by $Q_R = \{w^{-1}L | w^{-1}L \text{ is prime}\}$, $Q_{R0} = \{w^{-1}L \in Q_R | w^{-1}L \subseteq L\}$, $F_R = \{w^{-1}L | \varepsilon \in w^{-1}L\}$, and $\delta_R(w^{-1}L, a) = \{v^{-1}L \in Q_R | v^{-1}L \subseteq (wa)^{-1}L\}$.

\mathcal{R}_L is minimal with respect to the number of states (see [4], Theorem 1).

	e	e_1	e_2	e_3	e_4
s_1	1	0	1	1	0
s_2	1	1	0	1	1
s_3	1	0	1	0	0
s_4	0	0	0	0	1

Figure 1: An example for a coverable column (labeled by e)

3 Inferring a RFSA using an observation table

3.1 A “parasitic” two-step algorithm

The learner we define infers the canonical RFSA for L from a suitable combination of information sources. A source can be an oracle for membership queries (MQs; ‘Is this string contained in the language?’) or equivalence queries (EQs; ‘Is A a correct automaton for L ?’ – yielding some $c \in (L \setminus \mathcal{L}(A)) \cup (\mathcal{L}(A) \setminus L)$ in case of a negative answer) or a positive or negative sample of L fulfilling certain properties, and other kinds of sources can be considered as well. Suitable known combinations are: An oracle for MQs and EQs (a *minimally adequate teacher*, or MAT), an oracle for MQs with positive data, or positive and negative data.

In a first step we use an existing algorithm to build a table $T' = (\text{RED}' \cup \text{BLUE}', E', obs')$ representing the canonical DFA for the reversal \bar{L} of L . For eligible algorithms for various settings see [1] (L^* , MAT learning), [9] (learning from MQs and positive data), or [12] (this meta-algorithm covers MAT learning, MQs and positive data, and positive and negative data, and can be adapted to other combinations). All these learners add elements to the set labeling the rows of a table (candidates for states in \mathcal{A}_L) until it is closed, and/or separating contexts (i.e., suffixes revealing that two states should be distinct) to the set labeling the columns until it is consistent – additions of one kind potentially resulting in the necessity of the other and vice versa – and, once the table is closed and consistent, deriving a DFA from it that is either \mathcal{A}_L or can be rejected by a counterexample from the information sources, which is evaluated to restart the cycle. Obviously, since the sources only provide information about L and not \bar{L} , we must minimally interfere by adapting data and queries accordingly: Strings and automata have to be reversed before submitting them to an oracle, samples and counterexamples before using them to construct T' .

In the second step we submit T' to the following modifications:

- (1) Only keep one representative for every distinct row occurring in the table in RED' , and only keep one representative for every distinct column in E' .
- (2) Eliminate all representatives of rows and columns containing only 0s.
Let the resulting table be $T'' = (\text{RED}'' \cup \text{BLUE}'', E'', obs'')$.
- (3) Eliminate all representatives of *coverable* columns, i.e., all $e \in E''$ with

$$\begin{aligned} &\exists e_1, \dots, e_n \in E'' : \forall s \in \text{RED}'' : \\ &[obs''(s, e) = 0 \Rightarrow \forall i \in \{1, \dots, n\} : obs''(s, e_i) = 0] \wedge \\ &[obs''(s, e) = 1 \Rightarrow \exists i \in \{1, \dots, n\} : obs''(s, e_i) = 1]. \end{aligned}$$

For example, the column labeled by e in Figure 1 would be eliminated because its 1s are all “covered” by the columns labeled by e_1, e_2 , and e_3 .

Note that the first two modifications mainly serve to trim down the table to make the third modification less costly. In fact, most algorithms mentioned above can easily be remodeled such that they build tables

in which there are no rows or columns consisting of 0s and in which the elements labeling the rows in the RED part are pairwise obviously different already such that no row is represented twice.

The table thus modified shall be denoted by $T = (\text{RED} \cup \text{BLUE}, E, obs)$ and the derived automaton by $\mathcal{A}_T = (\Sigma, Q_T, Q_{T0}, F_T, \delta_T)$ with $F_T = F_{T'}$ (this has to be stated in case ε has been eliminated). As we have kept a representative for every distinct row and as all pairs of RED' elements that are distinguished by the contexts eliminated by (3) must be distinguished by at least one of the contexts covering those as well \mathcal{A}_T still represents $\mathcal{A}_{\overline{L}}$ (but without a failure state).

We use T to define $\mathcal{R} := (\Sigma, Q_R, Q_{R0}, F_R, \delta_R)$ with $Q_R = \{q \subseteq \text{RED} \mid \exists e \in E : s \in q \Leftrightarrow obs(s, e) = 1\}$, $Q_{R0} = \{q \in Q_R \mid \forall s \in q : obs'(s, \varepsilon) = 1\}$ (obs' in case ε has been eliminated), $F_R = \{q \in Q_R \mid \varepsilon \in q\}$, and $\delta_R(q_1, a) = \{q_2 \mid q_2 \subseteq \overline{\delta_T}(q_1, a)\}$ for $q_1, q_2 \in Q_R$ and $a \in \Sigma$, and $\overline{\delta_T}$ is the transition function of the reversal of \mathcal{A}_T .

Observe that every state in Q_R corresponds to a column in T . As every element of RED represents an equivalence class of \overline{L} under the Myhill-Nerode relation every state in Q_R also corresponds to a unique set of equivalence classes, and the associated column represents the characteristic function of that set.

We show that \mathcal{R} is the canonical RFSA for L . The proof uses Theorem 2.

Definition 8 Let $A = (\Sigma, Q, Q_0, F, \delta)$ be an NFA, and define $Q^\diamond := \{p \subseteq Q \mid \exists w \in \Sigma^* : \delta(Q_0, w) = p\}$. A state $q \in Q^\diamond$ is said to be coverable iff there exist $q_1, \dots, q_n \in Q^\diamond \setminus \{q\}$ for $n \geq 1$ such that $q = \bigcup_{i=1}^n q_i$.

Theorem 2 (Cited from [4]). Let L be regular and let $B = (\Sigma, Q_B, Q_{B0}, F_B, \delta_B)$ be an NFA such that \overline{B} is a RFSA recognizing \overline{L} whose states are all reachable. Then $C(B) = (\Sigma, Q_C, Q_{C0}, F_C, \delta_C)$ with $Q_C = \{p \in Q_B^\diamond \mid p \text{ is not coverable}\}$, $Q_{C0} = \{p \in Q_C \mid p \subseteq Q_{B0}\}$, $F_C = \{p \in Q_C \mid p \cap F_B \neq \emptyset\}$, and $\delta_C(p, a) = \{p' \in Q_C \mid p' \subseteq \delta_B(p, a)\}$ for $p \in Q_C$ and $a \in \Sigma$ is the canonical RFSA recognizing L .

As a further important result it has also been shown in [4], Section 5, that in a RFSA for some regular language L whose states are all reachable the non-coverable states correspond exactly to the prime RLs of L and that consequently Q_C can be identified with the set of states of the canonical RFSA for L .

Lemma 3 (See [4], Prop. 1). Let $A = (\Sigma, Q, Q_0, F, \delta)$ be a RFSA. For every prime RL $w^{-1}\mathcal{L}(A)$ there exists a state $q \in \delta(Q_0, w)$ such that $L_q = w^{-1}\mathcal{L}(A)$.

Theorem 4 \mathcal{R} is the canonical RFSA for L .

Proof. \mathcal{A}_T meets the conditions for \overline{B} in Theorem 2 as (a) all states of \mathcal{A}_T are reachable because \mathcal{A}_T contains no useless states, (b) \mathcal{A}_T is a RFSA: Every DFA without useless states is a RFSA (see [4]), and (c) $\mathcal{L}(\mathcal{A}_T) = \overline{L}$. As \mathcal{A}_T contains no useless states \mathcal{A}_T and $\overline{\mathcal{A}_T}$ have the same number of states and transitions, so we can set $B = \overline{\mathcal{A}_T} = (\Sigma, Q_T, F_T, Q_{T0}, \overline{\delta_T})$. Assuming for now that there is indeed a bijection between Q_R and Q_C it is rather trivial to see that

- there is a bijection between $Q_{R0} = \{q \in Q_R \mid \forall x \in q : obs'(x, \varepsilon) = 1\}$ and $Q_{C0} = \{p \in Q_C \mid p \subseteq F_T\}$ due to $F_T = \{x \in \text{RED} \mid obs'(x, \varepsilon) = 1\}$,
- there is a bijection between $F_R = \{q \in Q_R \mid \varepsilon \in q\}$ and $F_C = \{p \in Q_C \mid p \cap Q_{T0} \neq \emptyset\}$ due to the fact that $Q_{T0} = \{\varepsilon\}$, and that
- for every $q \in Q_R$, $p \in Q_C$, and $a \in \Sigma$ such that q is the image of p under the bijection between Q_R and Q_C , $\delta_R(q, a) = \{q_2 \in Q_R \mid q_2 \subseteq \overline{\delta_T}(q, a)\}$ is the image of $\delta_C(p, a) = \{p' \in Q_C \mid p' \subseteq \overline{\delta_T}(p, a)\}$.

It remains to show that there is a bijection between Q_R and the set of prime RLs of L , i.e., Q_C . From the definition of Q_R it is clear that \mathcal{R} is a RFSA: As noted above, every state in Q_R corresponds to a column in T , labeled by a context $e \in E$, and also to the set of equivalence classes $[s]_{\overline{L}}$ such that $se \in \overline{L}$ for $s \in \text{RED}$. As a consequence the reversal of the union of this set of equivalence classes equals the RL

$\bar{e}^{-1}L$, and hence every state in Q_R corresponds to exactly one RL of L . According to Lemma 3, there is a state in Q_R for each prime RL of L , so every prime RL of L is represented by exactly one column in T . By (3) we have eliminated the columns that are covered by other columns in the table. If a column is not coverable in the table the corresponding state in Q_R is not coverable either: Consider a column in the table which can be covered by a set of columns of which at least some do not occur in the table. Due to Lemma 3, these columns can only correspond to composed RLs of L . If we were to add representatives of these columns to the table they would have to be eliminated again directly because of the restrictions imposed by (3). This means that if a column is coverable at all it can always be covered completely by restricting oneself to columns that correspond to prime RLs of L as well, and these are all represented in the table. Therefore Q_R cannot contain any coverable states.

Thus the correspondence between Q_R and the set of prime RLs of L is one-to-one, and we have shown that \mathcal{R} is isomorphic to the canonical RFSA for L . \blacksquare

Corollary 5 *Let L be a regular language. The number of prime RLs of L is the minimal number of contexts needed to distinguish between the states of \mathcal{A}_L .*

Also note that we can skip the modification (3) in the second part of our algorithm if we restrict the target to bideterministic regular languages (see [2]).

3.2 Comparison to other algorithms: Query complexity

An advantage of the algorithm described above is the trivial fact that it benefits from any past, present, and future research on algorithms that infer minimal DFAs via observation tables, and at least until now there is a huge gap between the amount of research that has been done on algorithms inferring DFAs and the amount of research on algorithms inferring NFAs – or RFAS, for that matter.

A point of interest in connection with the concepts presented is the study of further kinds of information sources that could be used as input and in particular suitable combinations thereof (see for example [12] for a tentative discussion).

Another point of interest is complexity. As the second part of the algorithm consists of cheap comparisons of 0s and 1s only of which (3) is the most complex the determining factor is the complexity of the chosen underlying algorithm. One of the standard criteria for evaluating an algorithm is its time complexity, but depending on the different learning settings there are other measures that can be taken into consideration as well, one of which we will briefly address.

For algorithms that learn via queries a good criterion is the number of queries needed, obviously. The prototypical query learning algorithm, Angluin's [1] algorithm L^* , which can be seen in a slightly adapted version L_{col}^* in Figure 2, needs $O(I_L)$ equivalence queries and $O(|\Sigma| \cdot |c_0| \cdot I_L^2)$ membership queries, where I_L is the index of $L \subseteq \Sigma^*$ and $|c_0|$ the length of the longest given counterexample. By modifications the number of MQs can be improved to $O(|\Sigma|I_L^2 + I_L \log |c_0|)$ which according to [10] is optimal up to constant factors. On the other hand, it has been shown in [11] that it is possible to decrease the number of EQs to sublinearity at the price of increasing the number of MQs exponentially.

Recently, Bollig et al. [8] have presented a MAT algorithm for RFASs using an observation table that keeps very close to the deterministic variant L_{col}^* mentioned above. They introduce the notions of *RFSA-closedness* and *-consistency*.

Definition 9 *Let $T = (S, E, obs)$ be an observation table. A row labeled by $s \in S$ is coverable iff $\exists s_1, \dots, s_n \in S$ (is coverable by the rows of s_1, \dots, s_n iff)*

$$\begin{aligned} \forall e \in E : [obs(s, e) = 0 \Rightarrow \forall i \in \{1, \dots, n\} : obs(s_i, e) = 0] \wedge \\ [obs(s, e) = 1 \Rightarrow \exists i \in \{1, \dots, n\} : obs(s_i, e) = 1]. \end{aligned}$$

```

initialize  $T := (S, E, obs)$  with  $S = \text{RED} \cup \text{BLUE}$  and  $\text{BLUE} = \text{RED} \cdot \Sigma$ 
  by  $\text{RED} := \{\varepsilon\}$  and  $E := \{\varepsilon\}$ 
repeat until  $\text{EQ} = \text{yes}$ 
  while  $T$  is not closed and not consistent
    if  $T$  is not closed
      find  $s \in \text{BLUE}$  such that  $\text{row}(s) \notin \text{row}(\text{RED})$ 
       $\text{RED} := \text{RED} \cup \{s\}$  (and update the table via MQs)
    if  $T$  is not consistent
      find  $s_1, s_2 \in \text{RED}$ ,  $a \in \Sigma$ ,  $e \in E$  such that  $s_1a, s_2a \in S$ 
      and  $\neg(s_1 <> s_2)$  and  $\text{obs}(s_1ae) \neq \text{obs}(s_2ae)$ 
       $E := E \cup \{ae\}$  (and update the table via MQs)
    perform equivalence test
    if  $\text{EQ} = 0$  get counterexample  $c \in (L \setminus \mathcal{L}(\mathcal{A}_T)) \cup (\mathcal{L}(\mathcal{A}_T) \setminus L)$ 
       $E := E \cup \text{Suff}(c)$  (and update the table via MQs)
  return  $\mathcal{A}_T$ 

```

Figure 2: L_{col}^*

Let $\text{ncov}(S) \subseteq \text{row}(S)$ be the set of non-coverable rows labeled by elements in S .

Definition 10 Let $T = (S, E, obs)$ be an observation table. We say that a row $r \in \text{row}(S)$ includes another row $r' \in \text{row}(S)$, denoted by $r' \sqsubseteq r$, iff $\text{obs}(s', e) = 1 \Rightarrow \text{obs}(s, e) = 1$ for all $e \in E$ and $s, s' \in S$ with $\text{row}(s) = r$ and $\text{row}(s') = r'$.

Definition 11 A table $T = (\text{RED} \cup \text{BLUE}, E, obs)$ is RFSA-closed iff every row $r \in \text{row}(\text{BLUE})$ is coverable by some rows $r_1, \dots, r_n \in \text{ncov}(\text{RED})$.

Definition 12 A table $T = (\text{RED} \cup \text{BLUE}, E, obs)$ is RFSA-consistent iff $\text{row}(s_1) \sqsubseteq \text{row}(s_2)$ implies $\text{row}(s_1a) \sqsubseteq \text{row}(s_2a)$ for all $s_1, s_2 \in S$ and all $a \in \Sigma$.

From a RFSA-closed and -consistent table $T = (\text{RED} \cup \text{BLUE}, E, obs)$ Bollig et al. derive an NFA $\mathcal{R} = (\Sigma, Q_R, Q_{R0}, F_R, \delta_R)$ defined by $Q_R = \text{ncov}(\text{RED})$, $Q_{R0} = \{r \in Q_R \mid r \sqsubseteq \text{row}(\varepsilon)\}$, $F_R = \{r \in Q_R \mid \forall s \in \text{RED} : \text{row}(s) = r \Rightarrow \text{obs}(s, \varepsilon) = 1\}$, and $\delta_R(\text{row}(u), a) = \{r \in Q_R \mid r \sqsubseteq \text{row}(sa)\}$ with $\text{row}(s) \in Q_R$ and $a \in \Sigma$.

Theorem 6 (See [8]). Let T be a RFSA-closed and -consistent table and \mathcal{R}_T the NFA derived from T . Then \mathcal{R}_T is a canonical RFSA for the target language.

See [8] for the proof. The algorithm NL^* by Bollig et al. is given in Figure 3.

The theoretical query complexity of NL^* amounts to at most $O(I_L^2)$ EQs and $O(|\Sigma| \cdot |c_0| \cdot I_L^3)$ MQs. This exceeds the maximal number of queries needed by L_{col}^* in both cases which is due to the fact that with NL^* adding a context does not always lead to a direct increase of the number of states in the automaton derived from the table. Note that the authors of [8] show that their algorithm statistically outperforms L_{col}^* in practice, which is partly due to the fact that the canonical RFSA is often much smaller than the canonical DFA (see [4]). Nevertheless it is noteworthy that apparently inferring an automaton with potentially exponentially less states than the minimal DFA seems to be at least as complex.

Inspired by [8] we propose another parasitic two-step algorithm that uses an existing algorithm with access to a membership oracle to establish a table $T' = (\text{RED}' \cup \text{BLUE}', E', obs')$ representing \mathcal{A}_L and modifies it as follows:

```

initialize  $T := (S, E, obs)$  with  $S = \text{RED} \cup \text{BLUE}$  and  $\text{BLUE} = \text{RED} \cdot \Sigma$ 
  by  $\text{RED} := \{\varepsilon\}$  and  $E := \{\varepsilon\}$ 
repeat until  $\text{EQ} = \text{yes}$ 
  while  $T$  is not RFSA-closed and not RFSA-consistent
    if  $T$  is not RFSA-closed
      find  $s \in \text{BLUE}$  such that  $\text{row}(s) \in \text{ncov}(S) \setminus \text{ncov}(\text{RED})$ 
       $\text{RED} := \text{RED} \cup \{s\}$  (and update the table via MQs)
    if  $T$  is not RFSA-consistent
      find  $s \in S$ ,  $a \in \Sigma$ ,  $e \in E$  such that  $\text{obs}(sae) = 0$  and
         $\text{obs}(s'ae) = 1$  for some  $s' \in S$  with  $\text{row}(s') \sqsubseteq \text{row}(s)$ 
       $E := E \cup \{ae\}$  (and update the table via MQs)
    perform equivalence test
    if  $\text{EQ} = 0$  get counterexample  $c \in (L \setminus \mathcal{L}(\mathcal{A}_T)) \cup (\mathcal{L}(\mathcal{A}_T) \setminus L)$ 
       $E := E \cup \text{Suff}(c)$  (and update the table via MQs)
  return  $\mathcal{A}_T$ 

```

Figure 3: NL^* , the NFA (RFSA) version of L_{col}^*

(2)' Eliminate all representatives of rows and columns containing only 0s. Let $T'' = (\text{RED}'' \cup \text{BLUE}'', E'', \text{obs}'')$ be the resulting table.

(3)' For every $s \in \text{RED}''$ and every final state q_F of $\mathcal{A}_{T''}$ add an (arbitrary) string e to E'' such that $\delta_{T''}(\text{row}(s), e) = \{q_F\}$. Fill up the table via MQs.

Let $T = (\text{RED} \cup \text{BLUE}, E, \text{obs})$ be the resulting table. Note that as T' already contains the maximal number of possible distinct rows T is still closed and therefore RFSA-closed. T is RFSA-consistent as well: Recall that every element $s \in S$ represents a RL $s^{-1}L$ of L (see Section 2). If T was not RFSA-consistent we could find elements $s_1, s_2 \in S$, $e \in E$, and $a \in \Sigma$ with $\text{row}(s_1) \sqsubseteq \text{row}(s_2)$ but $\text{obs}(s_1a, e) = 1 \wedge \text{obs}(s_2a, e) = 0$. However, $ae \in s_1^{-1}L$ and $\text{row}(s_1) \sqsubseteq \text{row}(s_2)$ imply that $ae \in s_2^{-1}L$, and hence $\text{obs}(s_2a, e) = 0$ cannot be true.

From T we derive an automaton $\mathcal{R} = (\Sigma, Q_R, Q_{R0}, F_R, \delta_R)$ as in [8] (see above). The NFA \mathcal{R} is the canonical RFSA for L . This follows directly from Theorem 6 and the fact that T contains a representative for every RL of L .

The algorithm outlined above needs $I_L \cdot |F_L|$ MQs in addition to the queries needed by the algorithm establishing the original table but it does not require any more EQs. As EQs are usually deemed very expensive this can be counted in favor. Also note that if we restrict the target to bideterministic languages the table does not have to be modified and no additional queries have to be asked.

4 Conclusion

Two-step algorithms have the advantage of modularity: Their components can be exchanged and improved individually and therefore more easily adapted to different settings and inputs whereas non-modular algorithms are generally stuck with their parameters. One may doubt the efficiency of our two-step algorithms by observing that the second step partly destroys the work of the first, but as long as algorithms inferring the minimal DFA are much less complex than the ones inferring the minimal RFSA the two-step version outperforms the direct one.

It seems easy to adapt NL^* to other learning settings such as learning from positive data and a membership oracle or from positive and negative data in order to establish a more universal pattern for algorithms that infer a RFSA via an observation table similar to the generalization for DFAs attempted in [12].

References

- [1] Angluin, D.: Learning regular sets from queries and counterexamples. *Information and Computation* 75(2), 87–106 (1987)
- [2] Angluin, D.: Inference of reversible languages. *JACM*, vol. 29(3), pp. 741–765 (1982)
- [3] Hopcroft, J. E. and Ullmann, J. D.: *Introduction to Automata Theory, Languages, and Computation*. Addison-Wesley Longman (1990)
- [4] Denis, F., Lemay, A., and Terlutte, A.: Residual Finite State Automata. In: STACS 2001. LNCS, vol. 2010, pp. 147–155. Springer (2001)
- [5] Denis, F., Lemay, A., and Terlutte, A.: Learning regular languages using non-deterministic finite automata. In: ICGI. LNCS, vol. 1891, pp. 39–50. Springer (2000)
- [6] Denis, F., Lemay, A., and Terlutte, A.: Learning regular languages using RFSA. In: ALT 2001. LNCS, vol. 2225, pp. 348–363. Springer (2001)
- [7] Denis, F., Lemay, A., and Terlutte, A.: Some classes of regular languages identifiable in the limit from positive data. In: Grammatical Inference – Algorithms and Applications. LNCS, vol. 2484, pp. 269–273. Springer (2003)
- [8] Bollig, B., Habermehl, P., Kern, C., and Leucker, M.: Angluin-style learning of NFA. In: Online Proceedings of IJCAI 21 (2009).
- [9] Besombes, J. and Marion, J.-Y.: Learning Tree Languages from Positive Examples and Membership Queries. In: ALT. LNCS, vol. 3244, pp. 440–453. Springer (2003)
- [10] Balcazar, J.L., Diaz, J., Gavalda, R., and Watanabe, O.: Algorithms for learning finite automata from queries – a unified view. In: Advances in Algorithms, Languages, and Complexity, pp. 53–72 (1997)
- [11] Balcazar, J.L., Diaz, J., Gavalda, R., and Watanabe, O.: The query complexity of learning DFA. *New Generation Computing*, vol. 12(4), pp. 337–358. Springer (1994)
- [12] Kasprzak, A.: Meta-Algorithm GENMODEL: Generalizing over three learning settings using observation tables. Technical report 09-2, University of Trier (2009)
- [13] Kasprzak, A.: A learning algorithm for multi-dimensional trees, or: Learning beyond context-freeness. In A. Clark, F. Coste, L. Miclet (eds): ICGI 2008. LNAI, vol. 5278, pp. 111–124. Springer (2008)